Supercomputing and the Human Endeavor:
The Coming Scientific Revolution in
How We Use Machines to Help Us Think

Stephen M. Younger
Los Alamos National Laboratory
January 22, 2001

The world is about to undergo three major revolutions. Not revolutions in politics or government, but revolutions that will change the very fabric of our existence and the societies in which we live. Two are scientific, and one is social. The first, the revolution in the biological sciences, has been much discussed. The second, an imminent revolution in supercomputing and large scale simulation, is less well known. The third revolution is a social one that will be the direct result of the first two. In this paper we will discuss each of these revolutions, with a focus on the ethical and social implications of the revolution in supercomputing.

Why be concerned about a revolution in supercomputing? Because machines already under development will change the way that we think about the world and about ourselves. They will simultaneously create opportunities and dangers, just as have previous revolutions. Supercomputers will become the enabling engines of the knowledge economy, creating and transforming information at a fantastic rate. The access to such capabilities will exacerbate the growing "digital divide" between technological advanced countries and developing states. And, the application of supercomputing simulations to problems of psychology and biology will challenge humankind in a manner as profound as any philosophical or political revolution. It may even enable us to become the creators of our own successors. Something wonderful - and perhaps a little frightening - is about to happen.

Human beings today are remarkably similar to what they were at the dawn of recorded history, some five thousand years ago. We are a little taller than our ancestors because we eat more protein and we live quite a bit longer than they did because we have a better living environment and better health care. However, the three hundred generations that separate us from our most distant literate ancestors are not enough to allow the significant natural evolution of our bodies. The same is true on the intellectual side. We have made great progress in applying our brains to ever more difficult problems, but our basic mental capabilities have remained about the same over millennia. One can debate whether Aristotle or Einstein was the greater analytical intellect, but it is unarguable that genius has manifested itself throughout our known past.

Although revolutions have radically altered prevailing social institutions throughout the ages, our social systems have evolved at a more stately pace than one might naively think. The system of divine right monarchy that began in the ancient Near East persisted until the beginning of the twentieth century. Many of the core problems that human beings face today - how to raise children, how to run an equitable society that cares for the underprivileged, how to develop trust and trade among neighbors with different customs and language, and how to understand the world and our place in it - have been the focus of human study for thousands of years. Technology has improved the quality of our lives, but the fundamentals of daily life for most people have remained remarkably constant.

That is about to change. Advances in biology and computer science are about to present humankind with choices of an unprecedented nature, choices that touch upon what it means to be a human being, choices that touch upon where we wish our future as a species to lie. These advances build upon the intellectual revolution of the Greeks, in which the rules of rational thought was established. They build upon the revolution of Galileo in which analysis was combined with observation and experiment to yield what has become modern science. They build upon the industrial revolution in which this improved understanding was applied on a practical scale. And, they build upon the political revolutions that have swept the globe in the past several centuries, revolutions that have placed the critical decisions affecting society in the hands of those governed. We can and must learn from these previous revolutions in thought. We can and we must prepare the way for informed decisions to be made in the development and applications of new technologies that have broad implications for society and for the individual.

The revolution in biology, especially in the exploitation of the decoding of the human genome, has already been much discussed in scholarly and popular literature. For the first time in human history mankind will have the opportunity to change the human body in a significant way. Through the understanding of the human genome, the blueprint that contains all of the information required to make a human being, we will achieve the power to modify ourselves as well as our successors. We could choose to eliminate genetic diseases and extend our life span. We could contemplate the “improvement” of our species. Much has been written about the dangers of genetic engineering including the temptation to engineer a “better” human being. To many, this is a frightening prospect. There is something unique and precious about being human. Introducing changes that may have unintended and profound consequences is naturally worrying. Must we sacrifice on the altar of technological progress things in human nature that are beautiful and valuable? These are justifiable concerns. Not every opportunity that presents itself is a good one, nor is there any imperative that it be pursued. Genetic science is not a mandate - it is an opportunity, a choice that is presented to humanity for debate and consideration. There are wonderful benefits and dangerous pitfalls. Genetic engineering is one of the first choices that will affect us not as individuals or ethnic groups or nations but as an entire species. And, it is a choice that must be faced. The knowledge will come, the opportunities will arise. Human beings will be offered the choice to accept the natural course of development that we have followed for the history of life on earth or to leave the evolutionary track and become creators in our own right.

The second revolution, less discussed but equally profound, is the revolution in supercomputing and its potential for very large scale simulation. This revolution involves nothing less than changing the way that we use machines to help us think. Whereas the revolution in biology offers the opportunity to change the human body, the revolution in supercomputing offers the opportunity to change the human mind.

Several computers exist today that operate at speeds exceeding three trillion operations per second. A human being with a pencil and paper would take about one hundred thousand years to do the number of calculations that these computers can do in the time between two heartbeats. Put another way, each of these supercomputers is faster than several thousand of the fastest desktop computers. And the rate of development is breathtaking. In 2002 there will be a computer capable of thirty trillion operations per second. In 2004 peak speeds will reach one hundred trillion operations per second and by the end of the decade they will exceed one thousand trillion operations per second. These advances are more than “faster is better.” They will enable us to cross a threshold in how we interact with computers and what we use computers for. They will change the way that we use computers to help us think about ourselves and about the world around us. The change is qualitative, and is nothing short of revolutionary.

Progress in computing has been one of the greatest successes in human history. For the first five thousand years of human history, computation was limited by the speed of the human brain and our ability to keep track of numbers on fingers or abacus or paper. As late as 1940, the speed of general computation was still only about one operation per second. The advent of the electronic computer advanced computational speeds to hundreds of operations per second in the 1950’s, thousands in the 1960’s, millions in the 1990’s, and trillions in the first decade of the twenty-first century. Within a modern life span, the speed with which we can do calculations will have increased by a staggering 1,000,000,000,000,000 fold. Compare this to any other human achievement. In architecture we have progressed from buildings one story tall to skyscrapers of just over two hundred stories. In transportation we have advanced from a walking pace of about three miles per hour to the fastest spacecraft moving at a few tens of thousands of miles per hour. For many years computers have been able to best the most talented human savant in numerical computation. Recently, a computer beat the world chess champion. Even a simple projection of current technological trends leads one to the prospect that computers will soon cross a threshold in not only how fast they can solve problems, but in the type of problems that they can solve.

Until recently the computer has been just what its name implies - a device that performs numerical calculations. Such calculations could be performed just as effectively, if not as efficiently, by hand. Within the past twenty years computers have taken on a second important role, that of simulators of reality, with applications ranging from entertainment to the training of jet fighter pilots. Soon, supercomputers will move from performing calculations and simulating reality to addressing questions that touch upon the most basic issues of humanity.

Two examples will help to illustration this transition. At one thousand trillion operations per second, a speed that is expected to be achieved within the next decade, we will be able to think about simulating a living object at the atomic scale. This is more than a computational tour de force - it touches upon a question that human beings have been asking since human beings began to ask questions: What is life? What is the difference between living matter and dead matter? When does an organism cease to be a living one and become merely a collection of organic molecules? Conversely, what causes a collection of organic molecules to be alive?

At one thousand trillion operations per second, we will be able to construct a reasonably accurate model of at least a higher animal brain, including every synapse and neuron. We will be able to study brain function at a level of detail never before achieved. What was previously the domain of the humanist and psychologist will open to the computational scientist. Just as an understanding of the human genome will offer unprecedented insight into the origin of many physical diseases, large-scale simulation of brain function will offer an unprecedented insight into the origin of mental illness. Are mental disorders purely structural or chemical in origin or are some caused by individual experience in otherwise healthy brains, the collision of complex thought patterns? How might an improved understanding of brain function enable us to prevent and to treat mental illnesses? Freud dominated psychiatry during the first half of the twentieth century and psychopharmacology dominated treatment of many mental disorders during the second half of the century. Supercomputing may offer a third approach in which a greater precision is brought to psychotherapy. At the very least, the application of supercomputers to pharmacology will create a new and powerful tool for the design of drugs for chemically related mental disorders.

These examples are more than scientific “grand challenges.” Grand challenges are calculations we know that we could do if we only had a computer of sufficient speed and memory. Designing airplanes, modeling car crashes, making sense of complex observations from telescopes and particle accelerators are examples of grand challenges. The implications of calculating a living object or modeling a sophisticated brain go far beyond being able to answer a mathematical problem faster or more accurately. The answers to such “metaproblems” touch upon what it means to be human, what it means to live and to think and to appreciate beautiful things. Just as the scientific revolution in biology will bring fundamental choices related to our physical existence, the revolution in super-computing bring fundamental choices in what it means to be conscious human beings. Will understanding the operation of the human brain make us less human? Will such an understanding lead to the creation of horrible psychological weapons for warfare? Or, will it *emphasize* what it means to be human, distinguishing basic animal function from the wonder of humanity? What are the religious implications of understanding the detailed processes of life? Will these new computers offer the opportunity to create a fundamentally new type of intelligent life form? What would be our ethical respons-ibilities to such a life form? How will we explain its place in the universe, and how will we describe our own role in its creation?

There are many other questions that supercomputers may help us answer. At the same time, many more will arise. How do we think? What is the difference between the brain and the mind? How is the world that we perceive different from the world that actually exists? When does a “simulation” become “real”?

Knowing what is *not* possible is sometimes as important as knowing what *is* possible. What *cannot* be predicted or simulated using even a hypothetically *infinitely fast* computer? What are the *ultimate limits* of computer simulation? Already we know that long-range weather prediction is fundamentally impossible due to the chaotic nature of the atmosphere. No matter how powerful the computer, no matter how accurate the measurement of the initial conditions of the atmosphere, we now know that it is *fundamentally impossible* to predict accurate local weather conditions much more than thirty days in advance. Microscopically small variations in input conditions can grow exponentially in time to affect weather patterns thousands of miles away, as in the notion of tiny air currents produced by a butterfly in the Amazon jungle affecting the later course of a snowstorm in Siberia.

Another more personal limit to computation is the fundamental impossibility of duplicating a specific human brain. Human brains run on chemistry. And, chemistry is governed by the laws of quantum physics. Inherent in quantum physics is the principle of uncertainty that states that one can know only so much about a system and no more. At the atomic level there is a fundamental randomness to the universe, a randomness that causes events to happen with predictable probability but with unknowable certainty. One can say that two molecules have a statistical probability of reacting, but one cannot say with absolute certainty that two specific molecules will react at some precise time. The neurons and synapses in our brains are collections of complex molecules and they obey these same laws of quantum physics. Even if the exact structure and state of an individual human brain could be entered into a sufficiently powerful computer, the randomness inherent in quantum molecular processes would cause the computer simulation to diverge from the actual human brain. Individual human beings are unique and are destined by the most fundamental laws of the universe to remain unique. Even identical twins are not *absolutely* identical in their thought processes.

Can "consciousness" of any form reside in a computer? Can a computer have a "mind" as well as a "brain"? Some have argued that the fully deterministic and logical nature of computers precludes this possibility. Others wonder if a different *form* of consciousness might evolve, valuable in itself, but different from the consciousness of a human being. Fundamental considerations may limit the modeling of a specific human brain, but we may still be able to study the general principles of mental processes in a way that would be helpful in the treatment of mental disorders.

In these discussions we touch upon a fundamental issue in mathematics, the difference between was is *practically* impossible and what is *fundamentally* impossible. The equations that determine our weather appear at present to be subject to true chaos, a process that causes infinitesimal differences to grow exponentially. No matter how accurate our calculations, errors will creep and destroy their predictive accuracy. On the other hand, the equations of quantum molecular physics that govern the operation of the neurons and synapses in our brain do not appear to be chaotic. They are not fully deterministic in that quantum uncertainty applies, but we can be comforted to know that our mental processes do indeed appear to be logical and not chaotic!

An interesting aspect of the coming scientific revolution in supercomputing is the interface between humans and computers. Until now, we have dealt with computers on their terms. We type words on a keyboard, we manipulate a joystick or a mouse, we look at flat displays. Some progress has been made in three-dimensional imaging and in the use of gloves or other devices to couple body movements to computers. Contrast these means of interaction with how we interact with other human beings, with animals, or with the world in general. When we communicate with other people we speak in plain language. We "look people in the eye" to assess honesty. When we pet a dog we do so with hands that feel warmth and life. We see the world in three dimensions and integrate what we see with what we smell and hear and feel and taste. One of the spin-offs of the increased power of supercomputers is that we will be able to interact with them on *our* terms, in a more human fashion, using voice and vision and touch. However, bigger and faster is not necessarily *better* from a human perspective. A trivial illustration is the silliness of speed-reading poetry or listening to music at an accelerated speed. Even in mathematics we seek to know the elegant and simple underlying principles. Gigabytes of data generated by a computer may contain the answer, but we want to be able to understand and to comprehend and not just to calculate. We must appreciate *why* we are using the computer in order to understand *what* we want the computer to do for us.

The third revolution mentioned at the beginning of this essay is a social and geoeconomic revolution, related to how the world will come to grips with the two scientific revolutions. This third revolution has at least two major aspects: the effect on the individual and the effect on society. We have already discussed some of the impacts on the individual, including improvements in the treatment of physical and mental disease and potential changes in how we view ourselves as intelligent beings. The social implications of these revolutions are no less profound. The acceleration of progress in the biological and computer sciences will exacerbate differences in the distribution of wealth and opportunity among the developed and developing nations. Already the term "digital divide" is used to describe the gap between those who have access to computer technology and the Internet and those who lack such access. Supercomputing raises the issue of *access* to information to one of the ability to *create* new information. If the next economy is really a knowledge economy, supercomputers will play a vital role in the production of that knowledge. They will enable simulations to be performed more rapidly and more accurately, speeding time to market for new products. In the health arena they will enable the design of custom drugs for individuals, drugs that avoid the unwanted side effects of generic medications. The manufacturing of such drugs will require new technologies applied on a new scale, leading to growth in the industrial production sector.

Now is the time to begin serious dialog on the social and ethical implications of the revolution in supercomputing. This will join and extend the discussion that has been ongoing for decades in the biological sciences. Such dialog must include representatives of the scientific community to project from current technology what is doable in the future. It must also include historians to learn from past revolutions, philosophers to guide critical thought, religious thinkers to understand the implications for basic human values, economists to evaluate the effect on national wealth, sociologists to project social acceptance, and politicians to decide appropriate government investment. Such discussions should include those from the developed nations that will lead the revolution and those from the developing nations that contain much of the world's population, a population that can and must benefit from changes as profound as those discussed in this essay. Such changes have the opportunity to unite the world as never before in evaluating and choosing our future as human beings. It is an exciting prospect, but one that will not happen by itself.

The Woodrow Wilson Center and the Los Alamos National Laboratory have joined to sponsor a conference in June 2001 that will engage such a dialog among leading thinkers from around the world. *Supercomputing and the Human Endeavor* will provide a forum for identifying what is possible using projected technology and for discussing the implications for the individual and for society. Scientists and scholars, policy makers and economic leaders will have the opportunity to learn and to interact on a subject vital to the future of humanity.